IEEE Copyright Notice

arXiv:2008.04659v2 [eess.AS] 12 Dec 2021

S-vectors and TESA: Speaker Embeddings and a Speaker Authenticator Based on Transformer Encoder

N J Metilda Sagaya Mary, S Umesh, Sandesh V Katta

Abstract—One of the most popular speaker embeddings is x-vectors, which are obtained from an architecture that gradually builds a larger temporal context with layers. In this paper, we propose to derive speaker embeddings from Transformer's encoder trained for speaker classification. Self-attention, on which Transformer's encoder is built, attends to all the features over the entire utterance and might be more suitable in capturing the speaker characteristics in an utterance. We refer to the speaker embeddings obtained from the proposed speaker classification model as s-vectors to emphasize that they are obtained from an architecture that heavily relies on self-attention. Through experiments, we demonstrate that s-vectors perform better than x-vectors. In addition to the s-vectors, we also propose a new architecture based on Transformer's encoder for speaker verification as a replacement for speaker verification based on conventional probabilistic linear discriminant analysis (PLDA). This architecture is inspired by the next sentence prediction task of bidirectional encoder representations from Transformers (BERT), and we feed the s-vectors of two utterances to verify whether they belong to the same speaker. We name this architecture the Transformer encoder speaker authenticator (TESA). Our experiments show that the performance of s-vectors with TESA is better than s-vectors with conventional PLDA-based speaker verification.

Index Terms—s-vectors, speaker classification, speaker embeddings, speaker verification, TESA, Transformer encoder, x-vectors.

I. INTRODUCTION

SPEAKER verification uses speech as a biometric to verify the identity claimed by the speaker. There are two types of speaker verification systems: text-dependent and text-independent. Text-independent systems are flexible, as there is no constraint on the text spoken by the speaker. Most of the research in this area is focused on obtaining a single fixed-dimension vector representing an utterance. These vectors are then scored to verify the speaker's identity and are termed speaker embeddings. Any speaker embedding should enhance interspeaker variability and suppress intraspeaker variability while scoring.

One of the earliest speaker embeddings is i-vectors [1] extracted from the universal background model-Gaussian mixture model. With the increase in data available for training,

N J Metilda Sagaya Mary and S Umesh are with the Department of Electrical Engineering, Indian Institute of Technology Madras, Chennai 600036, India email:(metty789@gmail.com, umeshs@ee.iitm.ac.in)

Sandesh V Katta was with the Department of Electrical Engineering, Indian Institute of Technology Madras, Chennai 600036, India email:(sandesh.katta97@gmail.com)

speaker embeddings based on deep learning methods have gained in popularity.

A common approach in all the deep learning methods addressing the speaker verification task is to perform speaker classification first [2], [3], [4], [5], [6]. Then, utterance-specific fixed-dimension embeddings are obtained from the speaker classification network by different pooling methods. These embeddings are then fed to a speaker verification system to validate the identity claimed by the speaker. One such embedding is x-vectors [3] extracted from a speaker classification architecture based on time-delay neural network (TDNN) [7] with pooling. Hereafter, the term x-vectors refers to the original TDNN-based x-vectors in [3]. The x-vectors along with speaker verification based on probabilistic linear discriminant analysis (PLDA) [8] have significantly outperformed i-vectors.

Recently, speaker embeddings extracted from speaker classification architectures based on residual networks (ResNets) [9], [4], [5], [6], [10], [11] with pooling, which consider spectrogram representations of speech as images, have outperformed x-vectors. It should be noted that ResNets are pretrained models and are pretrained with a considerable number of images. These pretrained models are then retrained with speech data to perform speaker classification with standard softmax or additive margin softmax loss [12]. Then, speaker verification is performed with loss functions such as contrastive [13], [14] or relation loss [6].

Different strategies called pooling methods are employed in speaker classification networks to obtain fixed-dimension embeddings from variable length utterances, as mentioned before. These pooling methods can be classified into nontrainable and trainable approaches. Average [4] and statistics pooling [3] are nontrainable pooling methods. In average pooling, all the frames of the given dimension $D \times 1$ are averaged to obtain a single $D \times 1$ dimension vector. In statistics pooling, both mean and standard deviation vectors are concatenated to obtain a $2D \times 1$ vector. Attentive statistics [15], self-attention [16], [17], 2D self-attention [18], NetVLAD [19], and GhostVLAD [20] are trainable pooling methods. Not all frames are important in identifying a speaker. Therefore, attentive statistics and self-attention pooling aim to assign weights to the frames while calculating utterance level statistics. 2D self-attention pooling tries to incorporate an additional self-attention layer to pool the embeddings from different models. NetVLAD learns a $K \times D_k$ matrix, where K represents the number of clusters taken and D_k is the dimension of clusters. This matrix is then flattened and fed for speaker classification after some

processing. GhostVLAD is similar to NetVLAD, with a few clusters being dropped.

Recently, different attention mechanisms [21], [22] have been explored in natural language processing (NLP) and speech processing areas. A new architecture named Transformer [23] has outperformed many existing NLP models. Transformers have been successfully applied in both automatic speech recognition and text-to-speech tasks [24], [25], [26], [27]. A Transformer is built on self-attention across time. Some architectures for speaker verification have explored incorporating attention at a model level in frameworks different from Transformer. In [10], attention in both time and frequency was explored in a ResNet framework. In a TDNN framework, channelwise attention has shown improvement in speaker verification performance [28]. In [11], combining time-frequency attention with channel attention hierarchically at the model level was explored in a ResNet framework.

A Transformer has two modules: an encoder and a decoder. These two modules are built on self-attention and interact through source-attention. In the Transformer's encoder, information from all the frames is accounted for by the self-attention networks in every layer. Self-attention is built on dot products. As dot products are a kind of similarity measure that might better capture speaker characteristics in an utterance, we were motivated to explore embedding extraction from Transformer's encoder trained for the speaker classification task. It should be noted that this idea is different from attention-based pooling [15], [16], [17], [18], where attention is employed in the pooling only. At the time of submission, we also noticed a work that briefly discusses the extraction of embeddings from Transformer encoder [29]. [30] also explored the extraction of speaker embeddings from a Transformer encoder but with the objective of reducing the parameter count.

Different self-supervised pretrained models based on Transformers have given state-of-the-art performance in various NLP tasks. One such model is bidirectional encoder representations from Transformers (BERT) [31]. BERT is a Transformer encoder with some special input tokens and is trained on two tasks simultaneously: masked language modeling and next sentence prediction. Reconstruction of masked speech frames such as the masked words being predicted in the masked language modeling task of BERT was explored in [32], and the features extracted from it performed well in speaker classification as a downstream task. The next sentence prediction task of BERT takes two sentences as input, and the model is trained to predict whether the second sentence follows the first sentence. This next sentence prediction task motivated us to propose Transformer encoder speaker authenticator (TESA) in place of conventional PLDA.

Other works in speaker verification have been in the area of using different features. The impact of mel frequency cepstral coefficients (MFCCs) of different dimensions and fbank features in the x-vector framework was explored in [33]. Raw speech with a convolutional neural network-based speaker embedding extractor has also given the improvement in [33].

Our main contributions in this paper are as follows:
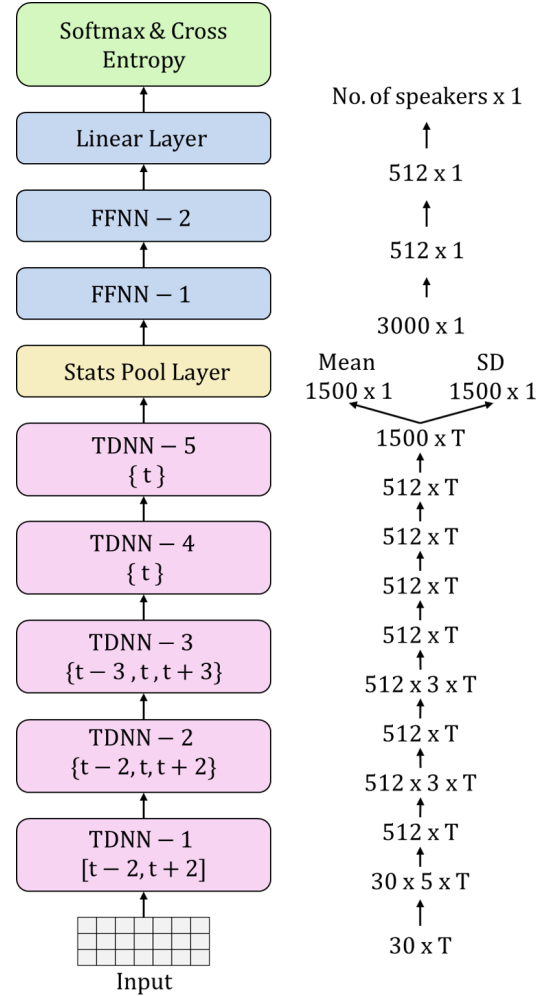- A Transformer encoder-based speaker classification archi-

Fig. 1: X-vectors architecture

tecture is proposed, from which we obtain better speaker embeddings. We call the embeddings obtained from this architecture s-vectors. We replaced the TDNN part of the x-vector model's architecture with Transformer's encoder to obtain the s-vector model's architecture without changing the input MFCC dimension.
- A Transformer encoder-based architecture for speaker verification is proposed, where we feed speaker embeddings from two utterances together, to verify whether they belong to the same speaker. We call this architecture the Transformer encoder speaker authenticator (TESA). TESA is inspired by the next sentence prediction task of BERT.
- Through experiments, we demonstrate that both the s-vectors and TESA outperform their existing baseline counterparts x-vectors and PLDA on both VoxCeleb-1 and VoxCeleb-2 data.

II. X-VECTORS

A. Architecture

The architecture of the x-vector model is shown in Fig. 1. We used 30-dimensional MFCC features instead of the 24-dimensional MFCC features in [3]. In the first TDNN layer

(TDNN-1), at any time step t, $\{t-2, t-1, t, t+1, t+2\}$ frames are spliced and presented as input. TDNN-2 takes $\{t-2, t, t+2\}$, and TDNN-3 takes $\{t-3, t, t+3\}$ spliced frames of the previous layer as input. TDNN-4 and TDNN-5 take just the t^{th} frame as input. To aggregate the statistics over the entire utterance, this system uses a statistics pooling layer (stats pool layer). This layer computes the mean and standard deviation across each dimension over the entire utterance, resulting in a single *3,000 × 1* vector. This single vector is representative of the entire utterance. Then, this vector is passed through two feedforward neural networks (FFNNs) and then to the output layer with cross-entropy as the criterion for speaker classification. All nonlinearities used are rectified linear units (ReLUs), and batch normalization is performed at every stage. The embeddings extracted from FFNN-1 before nonlinearity are termed x-vectors. This model is the baseline with which we compare our proposed s-vector model.

B. PLDA-based Speaker Verification

The x-vectors method uses PLDA to generate scores for the trial pairs. If the scores are higher than the set threshold, utterances in the pair belong to the same speaker and vice versa. Linear discriminant analysis helps verify seen speakers and cannot handle the unseen speaker scenario optimally. Therefore, PLDA attempts to fit a Gaussian mixture model with a continuous class variable to generalize for unseen speakers.

Let the probability of generating the data samples x from a given class y be given by

$$P(x|y) \sim \mathcal{N}(x|y, \phi_w), \tag{1}$$

where ϕ_w is the within-class covariance matrix and is common for all classes. To enable efficient handling of unseen classes, a Gaussian class prior is taken as

$$P(y) = \mathcal{N}(y|m, \phi_b), \tag{2}$$

where ϕ_b is the between-class covariance matrix and m is the mean of all data samples. ϕ_w is positive definite, and ϕ_b is positive semidefinite. A matrix V simultaneously diagonalizes both ϕ_w and ϕ_b as given by

$$\begin{aligned} V^T \phi_b V &= \psi \\ V^T \phi_w V &= I, \end{aligned} \tag{3}$$

where ψ is a diagonal matrix and I is the identity matrix. The class variable v and an example of a class u in the latent space are given by

$$\begin{aligned} v &\sim \mathcal{N}(.|0, \psi) \\ u &\sim \mathcal{N}(.|v, I). \end{aligned} \tag{4}$$

u and v are related to the feature space by

$$\begin{aligned} y &= m + Av \\ x &= m + Au, \end{aligned} \tag{5}$$

where A is V^{-T}.

Now, we have to test whether two examples (trial pair) from unseen classes belong to the same class. In the speaker verification task, classes are speakers, and examples are the
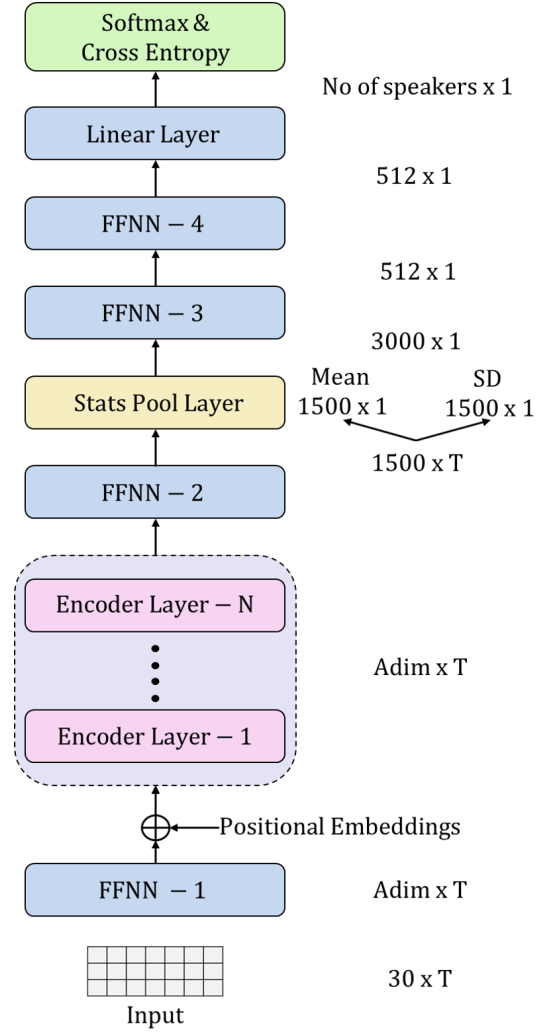

Fig. 2: S-vectors architecture

embeddings (x-vectors or proposed s-vectors) extracted from the utterances. To test whether the embeddings u_1 and u_2 of the two example utterances belong to the same speaker, the logarithm of the ratio of the likelihoods $P(u_1, u_2)$ and $P(u_1)$ $P(u_2)$ is computed. This is the PLDA score for the given trial pair, and if the score is greater than the set threshold, the utterances are taken to be from the same speaker and vice versa.

III. PROPOSED S-VECTORS

Existing speaker classification architectures gradually build temporal context with layers, as mentioned before. To capture the speaker characteristics better, we used the Transformer encoder, which is based on self-attention in our architecture. Its strength is that it is not restricted to finite context and attends to all frames in each of its layers as opposed to [16], [15], [17], [18], where attention is employed only during pooling.

A. Architecture

To derive s-vectors, we replaced the TDNN in the baseline x-vector model with the encoder of the Transformer [23],
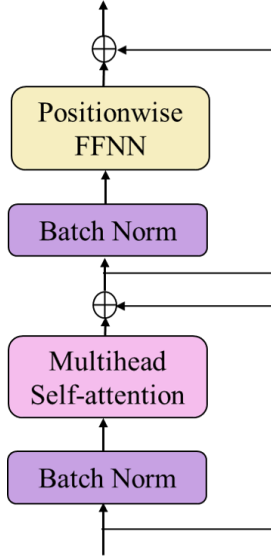
Fig. 3: Encoder layer architecture

$$U \in R^{T \times Adim}$$
$$W_Q^i, W_K^i \in R^{Adim \times d_k}$$
$$W_V^i \in R^{Adim \times d_v}$$
$$Q^i = U \times W_Q^i \quad (6)$$
$$K^i = U \times W_K^i$$
$$V^i = U \times W_V^i.$$

$O^i \in R^{T \times d_v}$ is the weighted average of the frames of V^i. The weights of the frames of V^i are obtained by,

$$O^i = \frac{softmax(Q^i K^{i^T})}{\sqrt{d_k}} V^i, \quad (7)$$

where softmax is computed along the row. In the case of multihead self-attention with P heads, the final resultant output (O) is given by

$$O = Concat(O^1, O^2, ..., O^p) \in R^{T \times (d_v \times P)}. \quad (8)$$

The value of d_k and d_v is set to $\frac{Adim}{P}$.

B. Datasets

We used two datasets in training our models: VoxCeleb-1 [34] and VoxCeleb-2 [6]. VoxCeleb-1 consists of recordings from 1,251 speakers and over 0.1 million utterances extracted from celebrity interview videos on YouTube. VoxCeleb-2 is approximately ten times larger than VoxCeleb-1. It consists of recordings from 6,112 speakers and has approximately a million utterances extracted from celebrity interview videos on YouTube. The minimum length of utterances in VoxCeleb-1 and VoxCeleb-2 is 3 seconds, and the maximum length is 20 seconds. The two datasets together amount to approximately 2,000 hours, and the male-female gender ratio is 61%-39%. The utterances are from speakers around the world belonging to different professions. Unless specified, all the models in this paper are trained on the whole of VoxCeleb-2 along with the Dev set of VoxCeleb-1. This training set is referred to as VoxCeleb-1+2 in this paper. Models are evaluated on VoxCeleb-1 Test. The VoxCeleb-1 Test consists of 4,874 utterances from 40 speakers and has 37,720 trial pairs.

C. Training Details

We used Kaldi's [35] VoxCeleb v2 recipe, which is the implementation of x-vectors in [3] with 30-dimensional MFCC features, to obtain our baseline. The same recipe was used to extract features and prepare data for s-vectors. After feature extraction, energy-based voice activity detection (VAD) is performed. The training data are then augmented with reverberation and noise in the same way as Snyder et al. for their x-vector model using VoxCeleb-1+2 data [3]. Reverberation examples are taken from the RIRS database [36], and noise examples are taken from the MUSAN database [37]. Cepstral mean normalization (CMN) is applied to the augmented utterances, and nonspeech frames are removed. They are then chunked to generate training examples, similar to the baseline x-vector model.

Our s-vector model was trained in ESPnet [38]. After data augmentation with RIRS and MUSAN, Kaldi's VoxCeleb v2

as shown in Fig. 2. Input is the same $30 \times T$ dimension MFCC features. We used the same training utterances as used by the x-vector system to ensure a fair comparison. $30 \times T$-dimensional MFCC features are transformed into the *attention dimension (Adim) $\times$ T* by FFNN-1 and fed to the above encoder layers after adding the position embeddings. Multihead self-attention is performed at every encoder layer. The encoder layer is explained in detail in the next section. The resultant *Adim $\times$ T* final encoder layer's output is then taken to *1,500 $\times$ T* through FFNN-2. Statistics pooling on these vectors results in a single *3,000 $\times$ 1* vector (*1,500 $\times$ 1* mean and *1,500 $\times$ 1* standard deviation). This *3,000 $\times$ 1* vector is then taken to *512 $\times$ 1* and then to *512 $\times$ 1* again by two FFNNs. The resultant vector is then presented to a classification layer. In all FFNNs except FFNN-2, we used ReLU. We used leaky ReLU (negative slope = 0.01) in FFNN-2 to stabilize the gradients flowing through the standard deviation of the stats pool layer. Speaker embeddings are extracted from the affine part of FFNN-3. We call the speaker embeddings extracted from our proposed model s-vectors.

1) Encoder Layer: Each encoder layer is made of a multihead self-attention network and a positionwise FFNN, as shown in Fig. 3. Positionwise FFNN is composed of two FFNNs. The first FFNN converts the input *Adim $\times$ T* to *encoder units $\times$ T*, and the second FFNN projects the output of the first FFNN back to *Adim $\times$ T*. Batch normalization is performed after adding residuals in every stage of the encoder layer. Batch normalization gave better performance than the usual layer normalization in the Transformers.

In the multihead self-attention network of the Transformer encoder, to obtain the output of a head i (O^i), the input (U) is converted to queries (Q^i), keys (K^i) and values (V^i) through the respective W_Q^i, W_K^i, W_V^i matrices as:

TABLE I: Training Details of S-vector Models Trained on VoxCeleb-1+2

Parameter	VoxCeleb-1+2
Position Encoding	Sinusoidal
Encoder Units	2,048
Normalize Before	True
Learning Rate	10
Batch Size	≈ 100
Optimizer	Noam
Dropout	0.1
Gradient-clip	5
Warm-up Steps	25,000
Normalization	Batch Norm

TABLE II: VoxCeleb-1+2 Results on VoxCeleb-1 Trials

Model	% EER	DCF (0.01)	DCF (0.001)
x-vector baseline	3.05	0.33	0.5
3L-256D-4H-S	3.35	0.34	0.53
6L-256D-4H-S	2.87	0.31	0.51
6L-512D-8H-S	**2.67**	**0.30**	**0.44**
9L-512D-8H-S	2.72	0.29	0.53

Note: 6L-512D-8H-S denotes a 6 encoder layer (L), 512 Adim (D), and 8 head (H) s-vector (S) model.
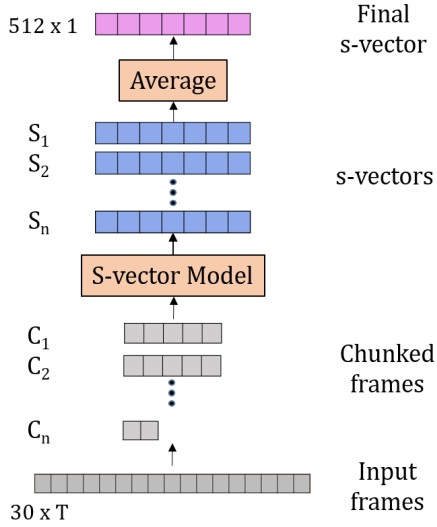
Fig. 4: Extraction of s-vectors by chunking

Fig. 5: Comparison of DET curves of the 6L-512D-8H-S model with the baseline x-vector model when trained on VoxCeleb-1+2 data
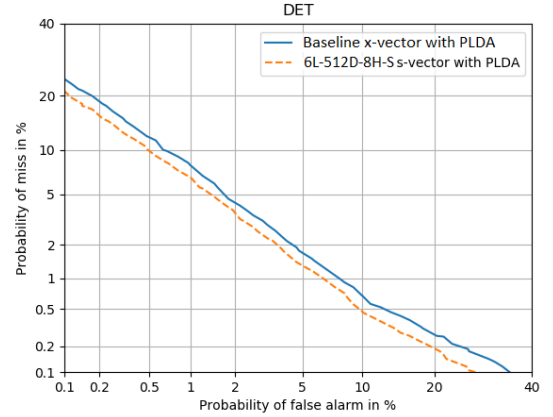
recipe chunks the data into random lengths with different start frames to generate training examples and writes it as .ark files before training. These ark files are not compatible with ESPnet. Therefore, we obtained chunk information, such as the start frame and chunk length, and performed chunking in ESPnet before passing an utterance for training. Therefore, the s-vector models were trained on the same data as that of the baseline x-vector model for a fair comparison. Other details of the s-vector training conducted in ESPnet are presented in Table I.

IV. Results with PLDA and Discussion

In this section, we discuss the results of the s-vector model with conventional PLDA scoring. We utilized Kaldi's VoxCeleb v2 recipe for PLDA scoring. Before feeding the utterances for s-vector extraction, CMN was performed, and nonspeech frames were removed similar to the x-vectors using an energy-based VAD system. Then, the resultant utterance frames were chunked, with each chunk being 300 frames, and the remaining frames were taken as another chunk. There was no overlap between the chunks. Embeddings were then extracted for each chunk and then averaged to obtain the s-vector, as shown in Fig 4. We tried different LDA dimensions and

found that the optimal dimension was 250 for the VoxCeleb-1+2 dataset.

A. Evaluation Metrics

The standard equal error rate (EER) and detection cost function (DCF) were used as the evaluation metrics to compare the baseline x-vector and our proposed s-vector model. EER refers to the value at which false alarm and miss error rates become equal. DCF is a weighted linear combination of false alarm and miss error rates. For DCF calculation, we assume P_{target} = 0.01 (or 0.001), while C_{miss} = 1 and $C_{false\ alarm}$ = 1.

B. Finding the Optimal Hyperparameters for the S-vector Model when Trained on VoxCeleb-1+2 Data

We analyzed the proposed s-vector architecture by varying the hyperparameters for VoxCeleb-1+2 data. % EER and DCF for the different numbers of layers and other hyperparameters taken are presented in Table II. 6L-512D-8H-S in Table II denotes a 6 encoder layer (L), 512 Adim (D), and 8 head (H) s-vector (S) model. The detection error tradeoff (DET) curves for the best performing s-vector and the baseline x-vector model are presented in Fig. 5. We see that the s-vector model outperforms the baseline x-vector model in terms of % EER by 12.5% relative, and DCF values are also better. The number of parameters in the 6L-512D-8H-S model is 25.3 million, and that of the baseline x-vector model is 4.3 million.

TABLE III: VOXCELEB-1 RESULTS ON VOXCELEB-1 TRIALS

Model	% EER	DCF (0.01)	DCF (0.001)
x-vector baseline	6.26	0.54	0.64
2L-256D-4H-S	5.80	0.48	0.65
3L-256D-4H-S	**5.63**	**0.46**	**0.62**
4L-256D-4H-S	5.50	0.50	0.64

Note: 3L-256D-4H-S means it is a 3 encoder layer (L), 256 Adim (D), and 4 head (H) s-vector (S) model.

TABLE IV: EFFECT OF CHUNKING ON 6-LAYER 512 DIMENSION 8 HEADS S-VECTOR MODEL TRAINED WITH VOXCELEB-1+2

Method	EER	DCF (0.01)	DCF (0.001)
6L-512D-8H-S-300	**2.67**	**0.30**	**0.44**
6L-512D-8H-S-100	2.96	0.35	0.53
6L-512D-8H-S-500	2.63	0.28	0.49

Note: 6L-512D-8H-S-300 means it is a 6 encoder layer (L), 512 Adim (D), and 8 head (H) s-vector (S) model with chunk length of 300.
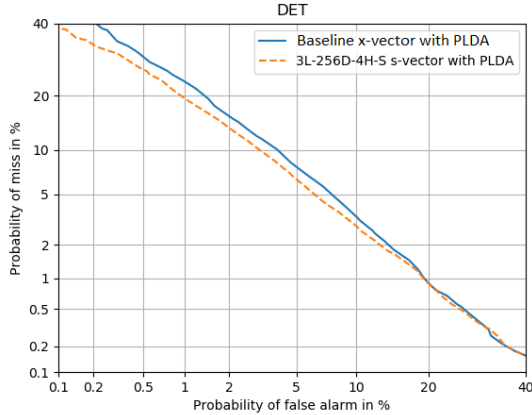

Fig. 6: Comparison of DET curves of 3L-256D-4H-S model with the baseline x-vector model when trained on VoxCeleb-1 data

Even the 6L-256D-4H-S model surpasses the x-vector baseline except for DCF (0.001) and has 12.8 million parameters.

C. Encoding Optimal Hyperparameters for the S-vector Model when Trained on VoxCeleb-1 Data

In this section, we show that the s-vector model outperforms the baseline x-vector model even with smaller training data. Both s-vector and the baseline x-vector models were trained with VoxCeleb-1 Dev data. These data are almost ten times smaller than the VoxCeleb-1+2 data. % EER and DCF for the different numbers of layers taken for the s-vector model are presented in Table III. We see that all three s-vector models consistently perform better than the baseline x-vector model in terms of % EER. The 3L-256D-4H-S model outperforms the baseline x-vector model trained on the same VoxCeleb-1 Dev data in both % EER and DCF. Therefore, the s-vectors outperform the x-vectors even in a dataset much smaller than VoxCeleb-1+2. The DET curves for the x-vector baseline and the best performing 3L-256D-4H-S model, when trained on the smaller VoxCeleb-1 Dev data, are presented in Fig. 6.

D. Effect of Chunking

In this section, we analyze the performance for different chunk lengths. We chose to chunk the utterance and obtain the embeddings for each utterance because of the position embeddings in our architecture. We expected that feeding whole utterances would lead to unseen positions in the input and result in poor speaker embeddings. Embeddings extracted for each chunk in an utterance are averaged to obtain the final

embedding. The performance of the 6L-512D-8H-S model for different chunk lengths is presented in Table IV. We see that the performance is better for chunk lengths of 300 and 500. A chunk length of 300 gives a good DCF (0.001) value, and we take it as the optimal chunk length.

E. Effect of Different Tapping Positions

We analyzed the effect of deriving embeddings from the affine part of FFNN-4 instead of FFNN-3. % EER and DCF for both of these positions of a 6L-512D-8H-S model trained on VoxCeleb-1+2 are presented in Table V. 6L-512D-8H-S-F3 means that the embeddings are tapped from the affine part of FFNN-3 of the 6L-512D-8H-S model. We see that FFNN-3 gives better embeddings than FFNN-4. Fig. 7 shows the tapping of embeddings from the affine part of FFNN-3 and FFNN-4.

TABLE V: EFFECT OF TAPPING LAYER ON 6L-512D-8H-S MODEL

Layer	EER	DCF (0.01)	DCF (0.001)
6L-512D-8H-S-F3	**2.67**	**0.30**	**0.44**
6L-512D-8H-S-F4	3.12	0.33	0.54

Note: 6L-512D-8H-S-F3 means it is a 6 encoder layer (L), 512 Adim (D), and 8 head (H) s-vector (S) model with embeddings extracted from FFNN-3 (F3).
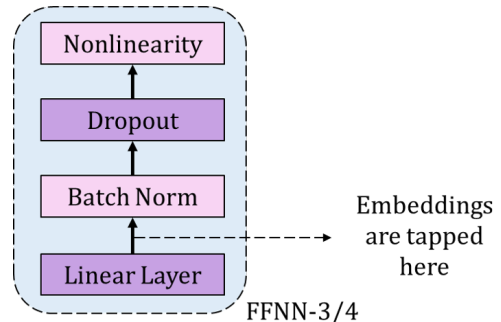

Fig. 7: Embedding tapping location in an FFNN

F. S-vector and X-vector Ensemble

As the s-vector and x-vector models are completely different in their fundamental architecture, we expected that the ensemble of these two models might improve the results. Therefore, the s-vectors of 6L-512D-8H-S and x-vectors from the
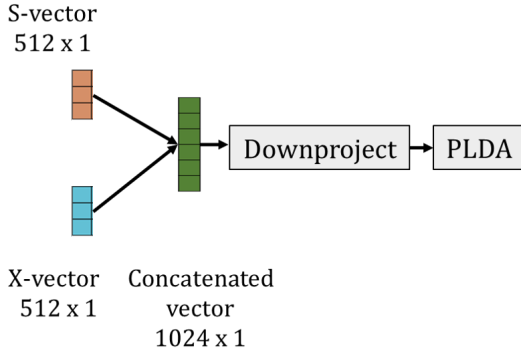
Fig. 8: S-vector and x-vector ensemble by concatenation

TABLE VI: PERFORMANCE OF S-VECTOR AND X-VECTOR ENSEMBLE

Method	EER	DCF (0.01)	DCF (0.001)
x-vector baseline	3.05	0.33	0.5
6L-512D-8H-S	2.67	0.30	0.44
Ensemble	**2.35**	**0.26**	**0.36**

baseline model were concatenated to obtain 1024-dimensional ensemble vectors. These vectors were then downprojected to 300 dimensions and then presented to PLDA for scoring, as shown in Fig. 8. As expected, the ensemble gives a significant improvement, as shown in Table VI.

V. PROPOSED TRANSFORMER ENCODER SPEAKER AUTHENTICATOR (TESA)

In PLDA, given a pair of enrollment and test embeddings, a log-likelihood score is computed to verify any identity claim. We propose a new speaker verification architecture, as a replacement for PLDA, in this section. This architecture is inspired by BERT.

A. TESA Architecture

BERT's next sentence prediction task is the inspiration for the Transformer encoder speaker authenticator (TESA). As mentioned before, the next sentence prediction task predicts whether the second text sentence in the sentence pair follows the first sentence. Therefore, we hypothesized that an architecture similar to BERT can be trained with embeddings from pairs of speech utterances to predict whether it is from the same speaker. In this case, the input is s-vector embeddings.

The data preparation for TESA is similar to the data preparation for BERT. However, we do not perform any masking because we are only interested in finding whether both utterances belong to the same speaker. It should also be noted that TESA operates on embeddings corresponding to a pair of spoken utterances and BERT operates on pairs of text sentences.

To train TESA, CMN was performed, and nonspeech frames were removed from every utterance. Then, s-vectors were extracted for every nonoverlapping 300 frame chunk of an utterance. Therefore, each utterance results in a set of embeddings, and every trial pair results in two sets of embeddings:
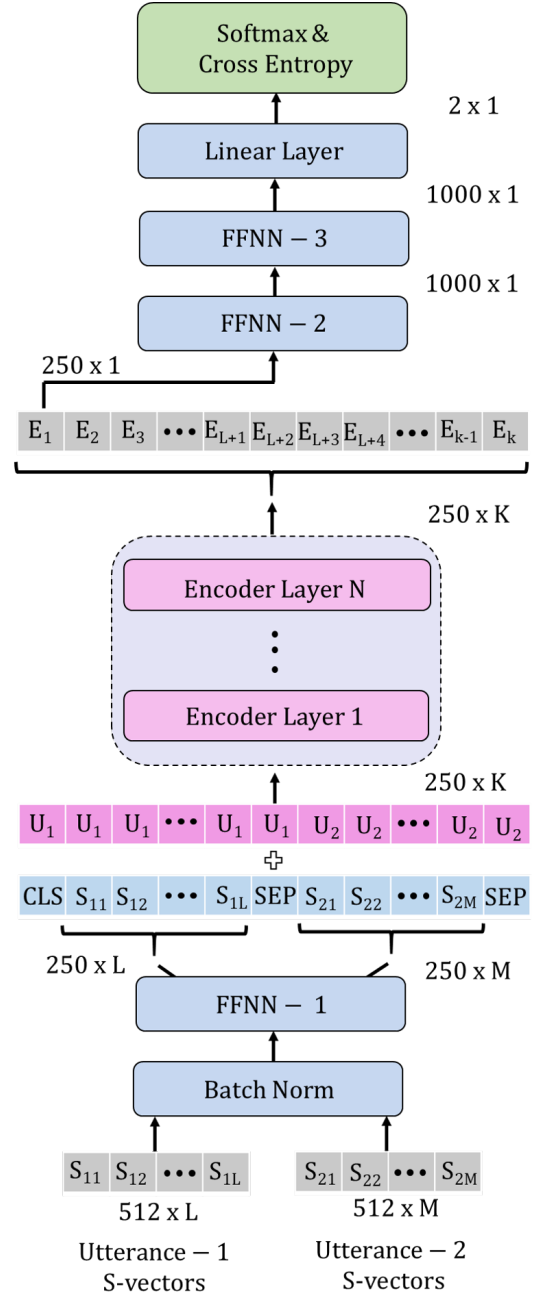

Fig. 9: Transformer Encoder Speaker Authenticator (TESA)

Utterance-1 s-vectors of dimension $512 \times L$ and Utterance-2 s-vectors with dimension $512 \times M$, where L and M are the number of chunks in the respective utterances.

Similar to BERT's data preparation, classifier (CLS) and separator (SEP) tokens are taken, as shown in Fig 9. Unlike BERT, SEP demarcates the boundary between the two utterances instead of sentences. The resultant has the dimension $512 \times K$, where $K = L + M + 3$.

The s-vectors are then added with learnable utterance embeddings (U_1, U_2) to indicate that the two sets of s-vectors belong to two different utterances. This is similar to adding sentence embeddings in BERT. We do not add position embeddings to TESA because the prediction should not depend
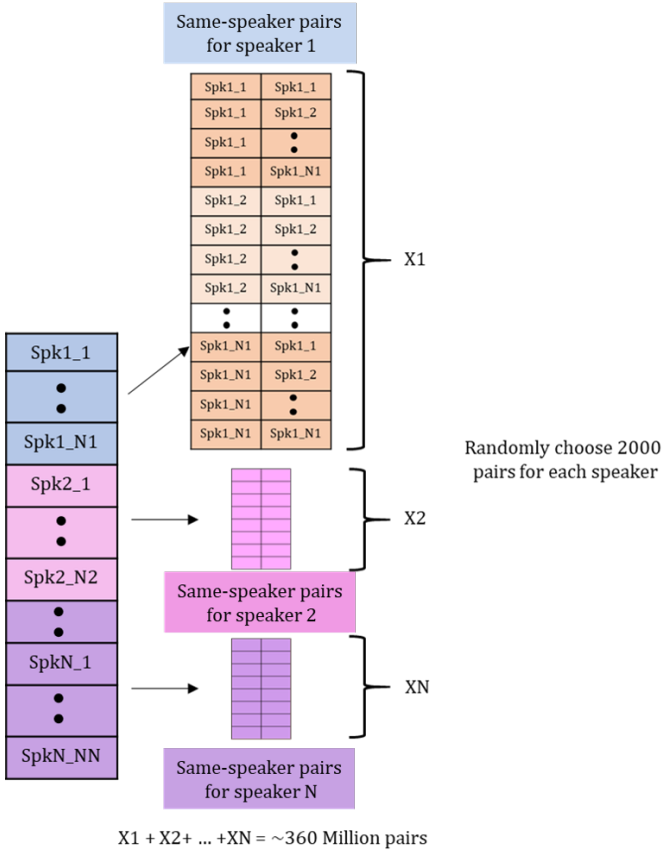
Fig. 10: Same-speaker dataset creation

Fig. 11: Different-speaker dataset creation for one speaker

TABLE VII: TRAINING DETAILS OF TESA

Hyperparameters	Details
Encoder Layers	9
Position Encoding	No
Encoder Units	1,024
Normalize Before	True
Learning Rate	10
Batch Size	2,000
Optimizer	Noam
Dropout	0.1
Gradient-clip	5
Warm-up Steps	25,000
Normalization	Batch Norm
Adim	250

on the order of the s-vector embeddings. The resultant is then fed to FFNN-1, which takes it to $250 \times K$. This is now fed to the Transformer encoder.

At the end of the Transformer encoder layers, E_1 corresponding to the input CLS token is tapped and fed to the next layers for classification, as shown in Fig. 9. Other outputs of the encoder ($E_2, E_3, ..., E_K$) are not considered. FFNN-2 takes the 250×1 vector to the $1,000 \times 1$ vector, and then FFNN-3 takes this $1,000 \times 1$ vector to another $1,000 \times 1$ vector. The linear layer projects the resultant vector to 2×1. Then, softmax is applied, and the model is trained on binary cross-entropy loss to predict whether the embeddings belong to the same speaker. We obtained the score for a given trial pair by subtracting the logits corresponding to the different-speaker classes from the same speaker class. All the nonlinearities used are ReLUs. We have performed batch normalization instead of layer normalization, similar to the s-vectors architecture.

B. Training Details

We used the same data PLDA uses for baseline x-vectors and extracted the s-vectors for chunks, as discussed in the previous subsection. The data have approximately 1.2 million utterances. After performing CMN and nonspeech frame removal, we created two sets of paired data: same-speaker and different-speaker. To create the same-speaker dataset, every utterance in the original dataset was paired with all possible same-speaker utterances, as shown in Fig 10. This resulted in
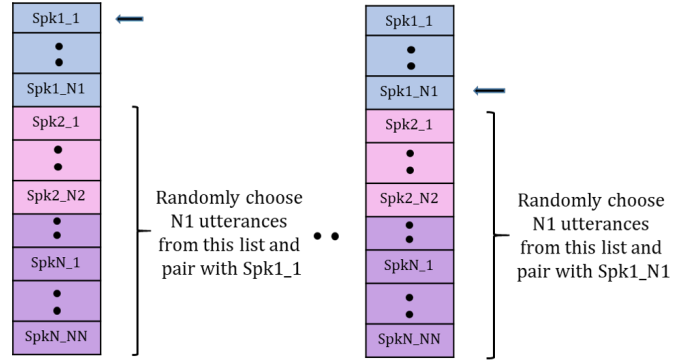
approximately 360 million pairs. To obtain 360 million pairs for the different-speaker dataset, every utterance of a particular speaker was paired with X randomly chosen different-speaker utterances, where X is the number of same-speaker utterances available, as shown in Fig 11. As the number of pairs was too high for the experiment, we randomly chose 2,000 pairs per speaker from both sets. This ensured the balance between the two classes. Some speakers resulted in fewer than 2,000 pairs due to fewer utterances pertaining to that speaker in the VoxCeleb dataset. For such speakers, all the available pairings were taken. These two datasets were then combined and shuffled to create the training dataset for TESA. All the experiments for TESA were performed in ESPnet. The hyperparameters are presented in Table VII.

VI. RESULTS WITH THE PROPOSED TESA AND DISCUSSION

All the results in section IV were obtained by PLDA scoring. In this section, we analyze the results of the Transformer encoder speaker authenticator (TESA) proposed in this paper. Table VIII shows the results obtained with the 9-layer TESA model. The number of parameters in TESA is 8.3 million. In Table VIII, it can be seen that TESA outperforms the conventional PLDA by a significant margin. We obtain a relative improvement of 11% in EER, and the DCF values are also better than those of conventional PLDA. DET curves for s-vectors with TESA, s-vectors with conventional PLDA
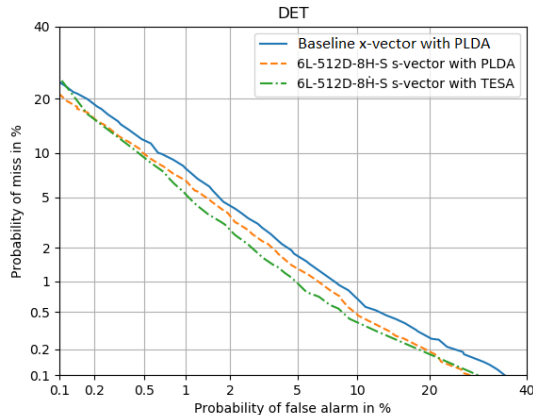
Fig. 12: Comparison of DET curves of s-vectors with TESA, s-vectors with PLDA and the baseline x-vector with PLDA

TABLE VIII: Comparison of TESA and PLDA for S-vectors from 6L-512D-8H-S

Encoder Layers	EER	DCF (0.01)	DCF (0.001)
s-vectors + PLDA	2.67	0.30	0.44
s-vectors + TESA	**2.37**	**0.28**	**0.39**

and the baseline x-vectors with conventional PLDA are shown in Fig. 12.

VII. CONCLUSIONS AND FUTURE WORK

In this work, we proposed deriving speaker embeddings from a speaker classification architecture based on Transformer's encoder. We call these embeddings s-vectors. The s-vectors obtained from our model trained on VoxCeleb-1+2 datasets outperformed the % EER of the baseline x-vector system trained on the same data, and gave a relative improvement of 12.5%. When trained only on the smaller VoxCeleb-1 dataset, the s-vectors again outperformed the % EER of the baseline x-vector system trained on the same data, and resulted in a relative improvement of 10%. Our model is also better in terms of DCF. In addition to s-vectors, we also proposed replacing conventional PLDA-based speaker verification with a new architecture named the Transformer encoder speaker authenticator (TESA). TESA outperformed the %EER of PLDA trained on the same dataset by 11% relative and has better DCF values. TESA with s-vectors jointly gives a relative improvement of 22.3% over the baseline x-vectors with PLDA on VoxCeleb-1 trials when trained on VoxCeleb-1+2 data. In the future, we would like to explore different pretraining methods and loss functions to improve the performance of the proposed architectures. We would also like to explore the capabilities of TESA trained in an end-to-end manner by feeding utterance pairs directly for speaker verification.

REFERENCES

[1] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front-end factor analysis for speaker verification," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 4, pp. 788–798, 2011.

[2] D. Snyder, D. Garcia Romero, D. Povey, and S. Khudanpur, "Deep neural network embeddings for text-independent speaker verification." in *Proc. of Interspeech*, 2017, pp. 999–1003.

[3] D. Snyder, D. Garcia Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *Proc. of International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5329–5333.

[4] J. S. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep speaker recognition," in *Proc. of Interspeech*, 2018, pp. 1086–1090.

[5] W. Xie, A. Nagrani, J. S. Chung, and A. Zisserman, "Utterance-level aggregation for speaker recognition in the wild," in *Proc. of International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5791–5795.

[6] A. Nagrani, J. S. Chung, W. Xie, and A. Zisserman, "VoxCeleb: Large-scale speaker verification in the wild," *Computer Speech & Language*, vol. 60, p. 101027, 2020.

[7] A. Waibel, T. Hanazawa, G. Hinton, K. Shikano, and K. J. Lang, "Phoneme recognition using time-delay neural networks," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 37, no. 3, pp. 328–339, 1989.

[8] S. Ioffe, "Probabilistic linear discriminant analysis," in *Proc. of European Conference on Computer Vision*. Springer, 2006, pp. 531–542.

[9] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. of Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2016, pp. 770–778.

[10] Y. Zhang, H. Yu, and Z. Ma, "Speaker verification system based on deformable CNN and time-frequency attention," in *Proc. of Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2020, pp. 1689–1692.

[11] C. Wang, J. Yi, J. Tao, Y. Bai, and Z. Tian, "Hierarchically attending time-frequency and channel features for improving speaker verification," in *Proc. of International Symposium on Chinese Spoken Language Processing (ISCSLP)*. IEEE, 2021, pp. 1–5.

[12] F. Wang, J. Cheng, W. Liu, and H. Liu, "Additive margin softmax for face verification," *IEEE Signal Processing Letters*, vol. 25, no. 7, pp. 926–930, 2018.

[13] S. Chopra, R. Hadsell, and Y. LeCun, "Learning a similarity metric discriminatively, with application to face verification," in *Proc. of Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2005, pp. 539–546.

[14] R. Hadsell, S. Chopra, and Y. LeCun, "Dimensionality reduction by learning an invariant mapping," in *Proc. of Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2006, pp. 1735–1742.

[15] K. Okabe, T. Koshinaka, and K. Shinoda, "Attentive statistics pooling for deep speaker embedding," in *Proc. of Interspeech*, 2018, pp. 2252–2256.

[16] Z. Lin, M. Feng, C. N. D. Santos, M. Yu, B. Xiang, B. Zhou, and Y. Bengio, "A structured self-attentive sentence embedding," *arXiv preprint arXiv:1703.03130*, 2017.

[17] Y. Zhu, T. Ko, D. Snyder, B. Mak, and D. Povey, "Self-attentive speaker embeddings for text-independent speaker verification." in *Proc. of Interspeech*, vol. 2018, 2018, pp. 3573–3577.

[18] G. Sun, C. Zhang, and P. C. Woodland, "Speaker diarisation using 2D self-attentive combination of embeddings," in *Proc. of International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5801–5805.

[19] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, "NetVLAD: CNN architecture for weakly supervised place recognition," in *Proc. of Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2016, pp. 5297–5307.

[20] Y. Zhong, R. Arandjelović, and A. Zisserman, "GhostVLAD for set-based face recognition," in *Proc. of Asian Conference on Computer Vision*. Springer, 2018, pp. 35–50.

[21] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," *arXiv preprint arXiv:1409.0473*, 2014.

[22] M. T. Luong, H. Pham, and C. D. Manning, "Effective approaches to attention-based neural machine translation," in *Proc. of Empirical Methods in Natural Language Processing*, 2015, pp. 1412–1421.

[23] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. of Advances in Neural Information Processing Systems (NIPS)*, 2017, pp. 5998–6008.

[24] L. Dong, S. Xu, and B. Xu, "Speech-transformer: A no-recurrence sequence-to-sequence model for speech recognition," in *Proc. of International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5884–5888.

[25] N. Li, S. Liu, Y. Liu, S. Zhao, and M. Liu, "Neural speech synthesis with transformer network," in *Proc. of AAAI Conference on Artificial Intelligence*, 2019, pp. 6706–6713.

[26] V. M. Shetty and N. J. Metilda Sagaya Mary, "Improving the performance of transformer based low resource speech recognition for Indian languages," in *Proc. of International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 8279–8283.

[27] N. J. Metilda Sagaya Mary, V. M. Shetty, and S. Umesh, "Investigation of methods to improve the recognition performance of Tamil-English code-switched data in transformer framework," in *Proc. of International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7889–7893.

[28] B. Desplanques, J. Thienpondt, and K. Demuynck, "ECAPA-TDNN: Emphasized channel attention, propagation and aggregation in TDNN based speaker verification," in *Proc. of Interspeech*, 2020, pp. 3830–3834.

[29] Y. Tong, W. Xue, S. Huang, L. Fan, C. Zhang, G. Ding, and X. He, "The JD AI speaker verification system for the FFSVC 2020 challenge." in *Proc. of Interspeech*, 2020, pp. 3476–3480.

[30] P. Safari, M. India, and J. Hernando, "Self-attention encoding and pooling for speaker recognition," *arXiv preprint arXiv:2008.01077*, 2020.

[31] J. Devlin, M. W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.

[32] A. T. Liu, S. W. Li, and H. y. Lee, "TERA: Self-supervised learning of transformer encoder representation for speech," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 2351–2366, 2021.

[33] W. W. Lin and M. W. Mak, "Wav2Spk: A simple DNN architecture for learning speaker embeddings from waveforms." in *Proc. of Interspeech*, 2020, pp. 3211–3215.

[34] A. Nagrani, J. S. Chung, and A. Zisserman, "VoxCeleb: A large-scale speaker identification dataset," in *Proc. of Interspeech*, 2017, pp. 2616–2620.

[35] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, "The Kaldi speech recognition toolkit," in *Workshop on Automatic Speech Recognition and Understanding*. IEEE Signal Processing Society, 2011.

[36] T. Ko, V. Peddinti, D. Povey, M. L. Seltzer, and S. Khudanpur, "A study on data augmentation of reverberant speech for robust speech recognition," in *Proc. of International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 5220–5224.

[37] D. Snyder, G. Chen, and D. Povey, "MUSAN: A music, speech, and noise corpus," *arXiv preprint arXiv:1510.08484*, 2015.

[38] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. E. Y. Soplin, J. Heymann, M. Wiesner, N. Chen *et al.*, "ESPnet: End-to-end speech processing toolkit," *arXiv preprint arXiv:1804.00015*, 2018.